
Subject: Re: [PATCH 12/25] sysfs: Introduce sysfs_rename_mutex

Posted by [ebiederm](#) on Wed, 08 Aug 2007 08:28:52 GMT

[View Forum Message](#) <> [Reply to Message](#)

Tejun Heo <htejun@gmail.com> writes:

> Hello, Eric.

>

> Eric W. Biederman wrote:

>> Looking carefully at the rename code we have a subtle dependency

>> that the structure of sysfs not change while we are performing

>> a rename. If the parent directory of the object we are renaming

>> changes while the rename is being performed nasty things could

>> happen when we go to release our locks.

>>

>> So introduce a sysfs_rename_mutex to prevent this highly

>> unlikely theoretical issue.

>

> Yeah, it's a theoretical issue. Rename/move implementation has always

> depended on the parent structure not changing beneath it, but it's nice

> to tighten up loose ends.

>

>> +DEFINE_MUTEX(sysfs_rename_mutex);

>

> Probably doesn't really matter but wouldn't a rwsem fit better?

Maybe. I didn't feel any loss in when I was writing the code.

Very few code paths actually seem to care.

>> @@ -774,7 +775,7 @@ static struct dentry *__sysfs_get_dentry(struct

> super_block *sb, struct sysfs_di

>> * down from there looking up dentry for each step.

>> *

>> * LOCKING:

>> - * Kernel thread context (may sleep)

>> + * mutex_lock(sysfs_rename_mutex)

Well this is weird in that it should be on sysfs_get_dentry

more than __sysfs_get_dentry but otherwise it's ok.

> LOCKING describes what locks should be held when entering the function,

> so proper description would be something like...

>

> Kernel thread context, grabs sysfs_rename_mutex

For rename_dir and move_dir yes. I was updating the rules

for sysfs_get_dentry. Which really wants it's parents to

hold that lock.

Eric

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
